

 Open access • Posted Content • DOI:10.1101/2021.05.26.445857

Expansion and contraction of resource allocation in sensory bottlenecks

— [Source link](#) 

Laura Rose Edmondson, Rodríguez Aj, Hannes P. Saal

Institutions: University of Sheffield

Published on: 27 May 2021 - bioRxiv (Cold Spring Harbor Laboratory)

Topics: Resource allocation and Bottleneck

Related papers:

- [Nonlinear scaling of resource allocation in sensory bottlenecks](#)
- [Functional diversity among sensory neurons from efficient coding principles](#)
- [Optimal Sensory Coding By Populations Of ON And OFF Neurons](#)
- [Metabolic cost of neuronal information in an empirical stimulus-response model](#)
- [Implicit embedding of prior probabilities in optimally efficient neural populations](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/expansion-and-contraction-of-resource-allocation-in-sensory-5ed3cm0tnd>

Expansion and contraction of resource allocation in sensory bottlenecks

Laura R. Edmondson^{1,3}, Alejandro Jiménez Rodríguez^{2,3}, and Hannes P. Saal^{1,3,✉}

¹Active Touch Laboratory, Department of Psychology, University of Sheffield, Sheffield S1 2LT, UK

²Department of Computer Science, University of Sheffield, Sheffield S1 3JD, UK

³Sheffield Robotics, University of Sheffield, Sheffield S1 3JD, UK

Topographic sensory representations often do not scale proportionally to the size of their input regions, with some expanded and others contracted. In vision, the foveal representation is magnified cortically, as are the fingertips in touch. What principles drive this allocation, and how should receptor density, e.g. the high innervation of the fovea or the fingertips, and stimulus statistics, e.g. the higher contact frequencies on the fingertips, contribute? Building on work in efficient coding, we address this problem using linear second-order models that maximize information transmission through decorrelation. We introduce a sensory bottleneck to impose constraints on resource allocation and derive the optimal neural allocation. We find that bottleneck width is a crucial factor in resource allocation, inducing either expansion or contraction. Both receptor density and stimulus statistics affect allocation and jointly determine convergence for wider bottlenecks. Furthermore, we show a close match between the predicted and empirical cortical allocations in a well-studied model system, the star-nosed mole. Overall, our results suggest that the strength of cortical magnification depends on resource limits.

efficient coding | decorrelation | cortical magnification | sensory representation | receptor density | stimulus statistics

Correspondence: h.saal@sheffield.ac.uk

Introduction

In many sensory systems, receptors are arranged spatially on a sensory sheet. The distribution of receptors is typically not uniform, but instead densities can vary considerably. For example in vision, cones are an order of magnitude more dense in the fovea than in the periphery (Goodchild et al., 1996; Wells-Gray et al., 2016). In the somatosensory system, mechanoreceptors are more dense in the fingertips than the rest of the hand (Johansson and Vallbo, 1979) (see also Fig. 1A). Alongside the density of receptors, the statistics of the input stimuli can also vary. For example, the fingertips are much more likely to make contact with objects than the palm (Gonzalez et al., 2014).

Subsequent sensory areas are often arranged topographically, such that neighbouring neurons map to nearby sensory input regions, for example, retinotopy in vision and somatotopy in touch. However, the size of individual cortical regions is often not proportional to the true physical size of the respective sensory input regions, and instead representations might expand (often called magnification) or contract (see also Fig. 1B). For example, both the fovea and the fingertips exhibit expanded representations in early visual and somatosensory cortices respectively, compared to their phys-

ical size (Azzopardi and Cowey, 1993; Engel et al., 1997; Martuzzi et al., 2014; Sereno et al., 1995). What determines such cortical magnification? For somatotopy, it has been proposed that cortical topography might directly reflect the density of sensory receptors (Catani, 2017). On the other hand, receptor density alone is a poor predictor of magnification (Corniani and Saal, 2020) and work on plasticity has established that cortical regions can expand and contract dynamically depending on their usage, suggesting that expansion and contraction might be driven by the statistics of the sensory stimuli themselves (Coq and Xerri, 1998; Merzenich and Jenkins, 1993; Xerri et al., 1996).

Here, we tackle this problem from a normative viewpoint, employing efficient coding theory, which has been widely used to model and predict sensory processing. Efficient coding theory (Barlow, 1961) suggests that neural populations are tuned to maximize the information present in the sensory input signals by removing redundant information (Atick, 1992; Atick and Redlich, 1990; Attneave, 1954; Chechik et al., 2006; Graham and Field, 2009). Efficient coding models have been most prominent in vision (Atick, 1992; Atick and Redlich, 1990; Bell and Sejnowski, 1997; Doi et al., 2012; Karklin and Simoncelli, 2011; Kersten, 1987; Olshausen and Field, 1996, 1997, 2004) and audition (Lewicki, 2002; Smith and Lewicki, 2006). This prior work has mostly focused on predicting the response properties and receptive field structure of individual neurons. In contrast, here we ask how receptive fields— independent of their precise structure—should tile the sensory input sheet when the inputs themselves differ in density and activation levels.

Some aspects of magnification in topographic representations have been qualitatively reproduced using self-organising maps (Ritter et al., 1992). However, these models generally lack a clear cost function and the magnification factor can be determined exactly only in rare cases, while a general expression is lacking (Ritter and Schulten, 1986). How varying receptor density should be taken into account has not been extensively studied, and most earlier work has assumed density of input receptors as constant (Atick, 1992; Doi et al., 2012) (but see Doi and Lewicki (2014) for an investigation of receptive field structure under different convergence ratios). However, signals from regions that are densely packed, such as the fovea, are likely to show greater spatial correlation than those from more sparsely innervated regions. These correlations should affect the resulting representation as efficient coding models aim to minimize redundancy. In contrast to re-

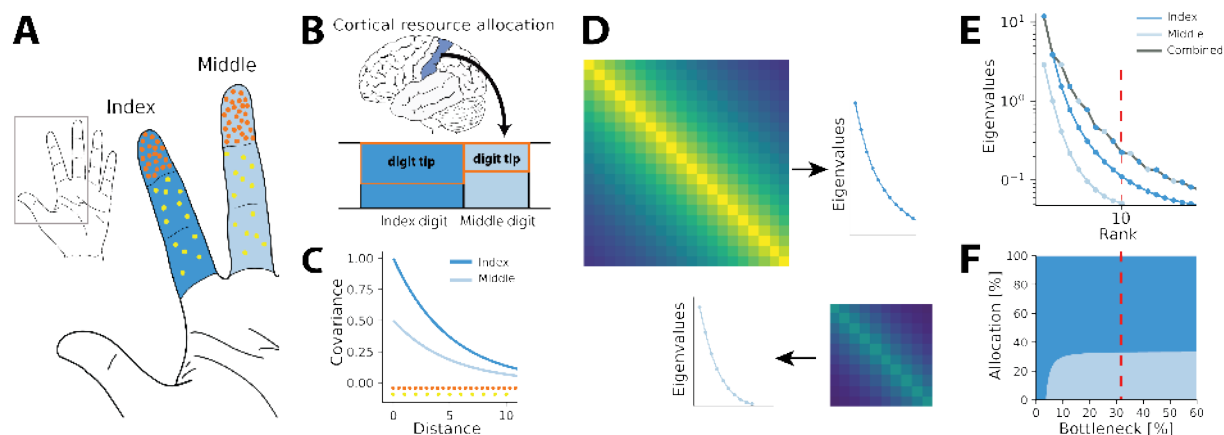


Fig. 1. Illustration of the resource allocation problem and solution outline. **A.** Example of problem setup. The density of touch receptors differs across regions of the hand (e.g. fingertips, shown in orange, versus finger base, yellow). Different fingers make contact with objects at different rates (dark blue versus light blue shading, darker colors indicating higher contact rates). **B.** How should cortical resources be divided between the different input regions? In the example shown here, the index digit could be allocated more resources than the middle digit due to increased activation. Similarly, the more densely innervated fingertips have larger allocation. **C.** Sensory inputs are correlated according to a covariance function (here, negative exponential) that decays with distance between receptors on the sensory sheet. This function is evaluated at different receptor distances depending on the density of sensory receptors (orange versus yellow dots at bottom). Regions with higher probability of activation exhibit greater variance (dark versus light blue curves). **D.** In efficient coding models that maximize decorrelation of sensory inputs in the presence of a bottleneck, the resource allocation problem can be solved by calculating the eigenvalues of the covariance matrix (defined by the covariance functions) for each region individually (here: example with two regions). **E.** The combined set of eigenvalues from all regions is then sorted; the region where each successive eigenvalue in the combined sorted set originates from determines where that output neuron's receptive field will fall. **F.** Counting which input regions successive eigenvalues belong to results in the allocation breakdown for different bottlenecks. For certain forms of the covariance function, this allocation can be calculated analytically.

ceptor density, there has been some work on how populations of neurons should encode non-uniform stimulus statistics using Fisher information (Ganguli and Simoncelli, 2010, 2014, 2016; Yerxa et al., 2020), an approximation of mutual information that can be used to calculate optimal encoding in a neural population (Yarrow et al., 2012). Rather than receptive fields uniformly tiling the input space, the optimal population should be heterogeneous, with receptive fields placed more densely over high probability inputs, at detriment to low probability regions (Ganguli and Simoncelli, 2014). Our approach differs from this prior work in two ways. First, rather than maximizing information between the neural population and the stimulus itself, we instead consider information between the neural population and an initial population of receptor neurons. This places a limit on the total amount of information that can be represented. Second, we consider the effects of limiting the population size, for which Fisher information is less suited, as it approximates mutual information best when population sizes are large.

The need for information maximization is often motivated by resource constraints. These can take the form of an explicit bottleneck, where the number of receptor neurons is greater than the number of output neurons. This is the case in the early visual system, where photoreceptors in the retina are much more numerous than the retinal ganglion cells to which they project (Goodchild et al., 1996; Wells-Gray et al., 2016). Other sensory systems might lack such explicit bottlenecks, but still place limits on the amount of information that is represented at a higher-order processing stage. For example, perceptual spatial acuity in touch is typically lower than might be expected from the spatial separation of mechanoreceptors, but can be improved with training up to a limit, suggesting that the relevant information is present in the sensory input but not typically represented in subsequent populations

(Wong et al., 2013). Sensory bottlenecks have been investigated in the context of efficient coding (Tishby et al., 2000; Tishby and Zaslavsky, 2015) and have had renewed theoretical interest, especially in deep learning (Lindsey et al., 2019; Schwartz-Ziv and Tishby, 2017; Tishby and Zaslavsky, 2015). How then should resource allocation change for different sized bottlenecks, given varying densities of receptors and different stimulus statistics? Here, we derive optimal neural allocations for different bottlenecks, while systematically varying receptor density and stimulus statistics. A preliminary version of these results restricted to the effects of receptor density in a one-dimensional space were previously presented as a conference paper (Edmondson et al., 2019).

Results

We consider linear second-order models which maximise information by decorrelating the sensory inputs. Decorrelation has featured prominently in models explaining early sensory processing, particularly in vision (Graham et al., 2006). Here, we use a decorrelation model to investigate the contribution of non-uniform receptor densities and activation to the resulting sensory representations (see Methods for details). Receptors can be placed in non-uniform densities across the sensory sheet, with some regions having a higher density than others. For example, in Fig. 1A, both digit tips are tiled more densely than the rest of finger. The receptor activation covariance decays with distance (see Fig. 1C). Thus, directly neighbouring units in denser regions covary more strongly than those from less dense regions. The activation of receptors can also vary between regions, which is modelled through scaling the covariance function (see appendix A for further details). We approximate the regions as separate, such that no receptors are activated from both regions simultaneously. The covariance

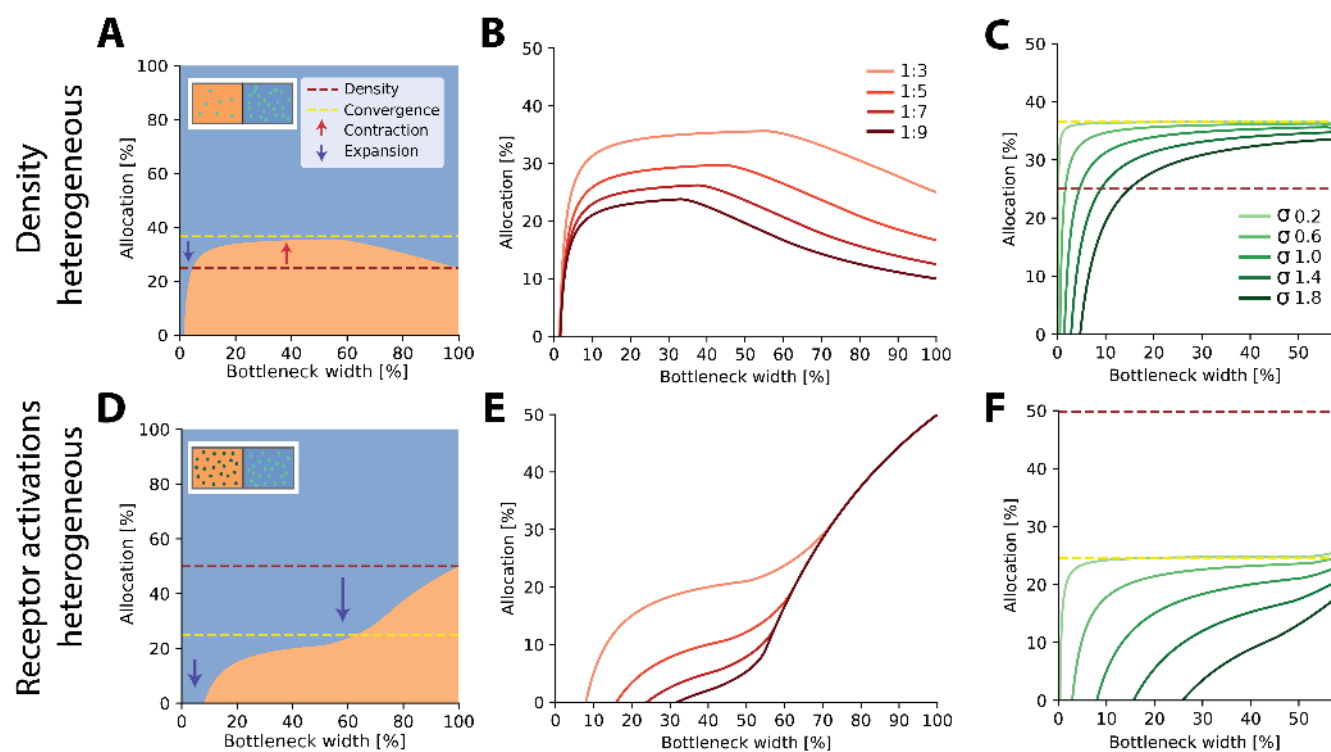


Fig. 2. Optimal resource allocation for heterogeneous receptor densities or input statistics leads to complex trade-offs. **A.** Illustration of resource allocation for heterogeneous receptor density but homogeneous stimulus statistics over all bottleneck sizes. Orange denotes the lower density region and blue the higher density region, with a ratio of 1:3. Dotted lines show proportional representation according to receptor numbers (red) and convergence of the optimal allocation in the limit (yellow). Arrows indicate contraction (up) and expansion (down) of the higher density region representation. **B.** Bottleneck allocation boundaries for different density ratios (given as low:high). The area below each line corresponds to the low density representation, while the area above corresponds to the high density representation, as in panel **A**. **C.** Effect of changing the extent of the spatial correlations (parameterized by the decay value σ , see Methods for details and Fig. S1 for an illustration of the covariance function for different values of σ). Density ratio is set at 1:3 for all σ . Increasing σ leads to expansion of the higher density region for a larger initial portion of the bottleneck. **D-F.** Same as in row above but for homogeneous density and heterogeneous receptor activation. **D.** Illustrative example with the blue region having higher receptor activation. Note that the representation of the higher activation region is expanded for all bottleneck widths. **E.** Allocation boundaries for different activation ratios. The representation of the high-activation regions is expanded for all bottlenecks. As activation ratio increases, the highly active region allocation is expanded for wider bottlenecks. **F.** Changing the extent of spatial correlations (σ) has larger effects when the activation is heterogeneous (set at 1:3 for all σ) compared to heterogeneous density (panel **C**). See Fig. S2 for an equivalent figure considering one-dimensional receptor arrangements.

matrix across all receptors then forms a block matrix, where the covariance between separate regions is zero (Fig. 1D). Eigenvalues can be calculated separately from the covariance matrices of each region (inset panels in Fig. 1D). These are then ordered by magnitude for both regions combined (grey line in Fig. 1E). A bottleneck is introduced by restricting the total number of output neurons. This is done by selecting from the combined set of ordered eigenvalues until the limit is reached. The proportion of eigenvalues originating from each input region determines its allocation for the chosen bottleneck width (see red dashed line Fig. 1E,F).

In the following, we focus on negative exponential covariance functions, for which the eigenvalues and the resulting allocation can be calculated analytically (see Methods for mathematical derivations). We will present results for the common case of two-dimensional sensory sheets (relevant for vision and touch, for example), while results for one-dimensional receptor arrangements are summarized in the appendix. For ease of analysis, all cases discussed will assume two sensory input regions of equal size, but differing in receptor density, input statistics, or both. Results for more than two regions or disparate region sizes are straightforward to calculate.

Resource limits determine the amount of magnification.

First, we investigated resource allocation in bottlenecks for heterogeneous density of receptors and heterogeneous stimulus statistics separately, whilst keeping the other factor uniform across the input regions.

Heterogeneous density.

For two regions with different receptor densities, we found that the higher density region could either expand or contract relative to its input density, depending on the width of the bottleneck. Specifically, for narrow bottlenecks (smaller than approximately 10% of input neurons), the higher density region is either exclusively represented or its representation is expanded compared to a proportional density allocation (see example in Fig. 2A). Mathematically, this can be explained by a multiplicative scaling of the eigenvalue function for the higher density region (see illustration in Fig. S3A). In contrast, for intermediate bottlenecks, low density regions expand their representation, beyond what is expected for a proportional mapping (see dashed red line in Fig. 2A denoting proportional allocation), leading to a contraction of the high density region. For negative exponential covariance

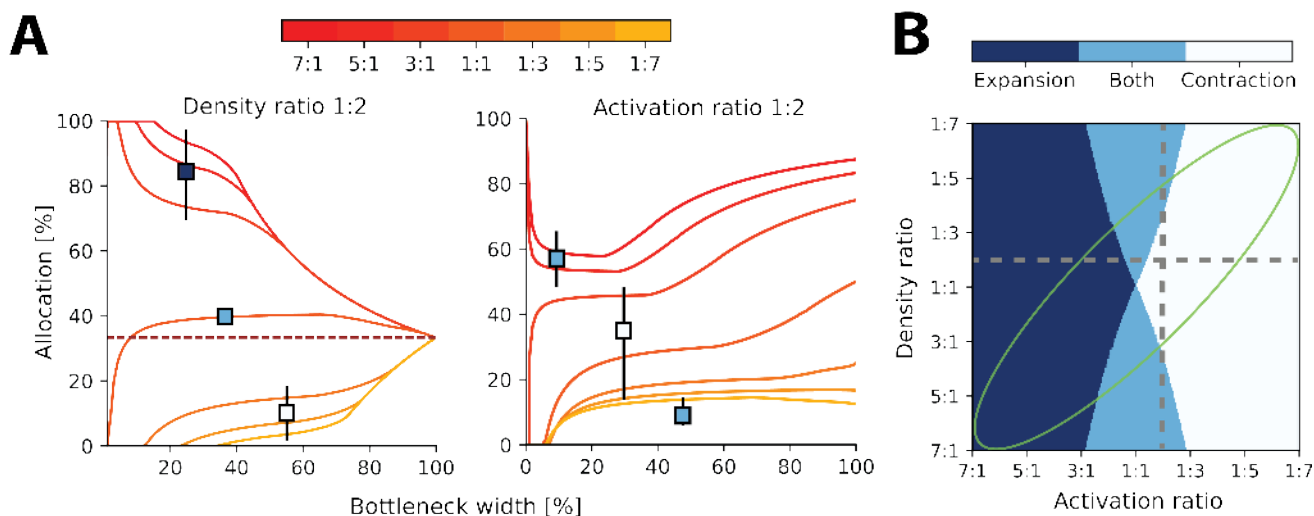


Fig. 3. Interactions between heterogeneous statistics and density. **A.** Expansion and contraction for a region with constant density and activation (baseline), while both density and activation are varied for the other region. All ratios given are baseline:other region. Symbols: *Dark blue Square* denotes expansion of the baseline region. Box lines show all combinations with expansion. *Light blue square*, expansion and contraction of the baseline region. *White square* indicates contraction of the baseline region. Left: Fixed density ratio, while activation ratio is varied. Right: Fixed activation ratio, while density ratio is varied. **B.** Possible expansion/contraction regimes for baseline region based on combinations of density and activation ratios. Grey dotted lines show all possible allocation regimes for a region if the activation ratio (vertical) or density (horizontal) is fixed. The green ellipse highlights parameter combinations where activation and density are correlated. See Fig. S4 for a comparison of how receptor density and activation interact between one-dimensional and two-dimensional receptor arrangements.

functions, this allocation converges to a fixed ratio at wider bottlenecks of $1/(1+\sqrt{a})$, where d is the density ratio between both regions (dashed yellow line in Fig. 2A; see Methods for derivation), as neurons are allocated to either region at a fixed ratio (see inset in Fig. S3A). Finally, for wide bottlenecks all information arising from the low density region has now been captured, and any additional output neurons will therefore be allocated to the high density region only. The over-representation of the low density region thus decays back to the original density ratio. The overall nonlinear effect of bottleneck width is present regardless of the ratio between the densities (see Fig. 2B). The spatial extent of the correlations over the sensory sheet (controlled by the decay constant σ in the covariance function, see Methods) determines allocation at narrow bottlenecks, and how fast the allocation converges, but does not affect the convergence limit itself (Fig. 2C). As σ increases, and therefore spatial correlations decrease, the convergence limit is approached only at increasingly wider bottlenecks. The extent of magnification thus heavily depends on the correlational structure of the stimuli for narrower bottlenecks, while receptor densities are more important for wider bottlenecks.

Heterogeneous statistics.

Aside from receptor densities, the stimulus statistics can also vary over the input space, leading to differences in receptor activation levels and their associated response variance (see appendix A for further details). Overall, allocations for heterogeneous receptor activation are similar to those found with heterogeneous density. However, while the allocations are again a nonlinear function of bottleneck width, the representations are solely expanded for the high activation region across the entire bottleneck (see example in Fig. 2D). The extent of this expansion depends on the width of the bottleneck,

and is again more extreme for narrower bottlenecks (see examples in Fig. 2E). The convergence limit is $1/(1+a)$ where a is the activation variance ratio, meaning that the level of expansion and contraction in intermediate bottlenecks is more extreme than in the heterogeneous density case (see difference between red and yellow dashed lines in (Fig. 2A and D). Finally, the effect of spatial correlations is also more pronounced (Fig. 2F).

In the cases described above, the bottleneck was constrained by the number of output neurons. Alternatively, the limiting factor might be set as the amount of information (total variance) captured. Doing so leads to similar allocation curves that retain the nonlinear behaviour described here (see Fig. S5 for some illustrations).

In summary, we find that representations of different inputs regions can contract or expand, depending on the bottleneck width. This effect plays out similarly for differences in receptor density and receptor activation, however with some crucial differences. Finally, for narrow bottlenecks the spatial extent of the correlations across the sensory sheet becomes an important driver.

Interplay between stimulus statistics and receptor density.

In sensory systems, such as in touch, both the density and input statistics vary across regions and will therefore jointly determine the resulting allocation. As a consequence, the convergence for intermediate bottlenecks will depend on both density and activation ratios, and can be calculated as $1/(1+a\sqrt{d})$, where a is the activation and d is the density ratio (see Methods for derivation). Importantly, for two-dimensional sensory sheets the activation ratio carries higher weight in determining the allocation than does the density ratio, and this effect is more pronounced at narrower bottle-

necks (see much wider spread of possible allocations when varying activation in the left panel of Fig. 3A than when varying density as shown in the right panel). Specifically, this means that the allocation regime, i.e. whether the allocation expands, contracts or exhibits both behaviours across all bottleneck widths, is more dependent on receptor activation than densities (Fig. 3B).

Finally, it is likely that regions with higher receptor densities will also show greater activation than lower density regions. For example, in touch, the regions on the hand with the highest receptor densities also are the most likely to make contact with objects (Gonzalez et al., 2014). In these cases, both effects reinforce each other and drive the resulting allocation further from proportional allocation (see orange lines in Fig. 3A, and green ellipse in Fig. 3B).

Resource limits determine the strength of plasticity under changes in stimulus statistics.

Over most of the lifetime of an organism the resources available for processing sensory information and the density of sensory receptors should be relatively constant. The stimulus statistics, on the other hand, can and will change, for example when encountering a new environment or learning new skills. These changes in stimulus statistics should then affect sensory representations, mediated by a variety of plasticity mechanisms. For example, increased stimulation of a digit will lead to an expansion of that digit's representation in somatosensory cortex (Jenkins et al., 1990).

We asked how representations should adapt under the efficient coding framework, and whether resource limits would affect the resulting changes. To answer this question, we calculated optimal allocations for different bottleneck widths, receptor densities, and stimulus statistics. We then introduced a change in stimulus statistics and re-calculated the resulting allocations (see illustration in Figure 4A). As expected, we found that when increasing the receptor activation over a region (for example, by increasing stimulation of the region), more neurons would be allocated to this region. Interestingly, however, this effect was dependent on the width of the bottleneck. The largest effects are seen for smaller bottlenecks, and then diminish as the bottleneck size increases. Figure 4B demonstrates such allocation changes for three different bottleneck widths for a range of receptor densities and activation ratios. This suggests that plasticity should be relatively stronger under extreme resource constraints than in cases where limits on the information are generous.

Predicting cortical magnification in the star-nosed mole.

Finally, we investigated to what extent the procedure outlined in the previous sections might predict actual resource allocation in the brain. While our model is relatively simple (see Discussion), decorrelation has been found to be a strong driver in early sensory processing and one might therefore expect the approach to at least yield qualitatively valid predictions. As currently available empirical data makes it difficult to test the impact of different bottleneck widths on the resulting allocation directly (see Discussion), we instead fo-

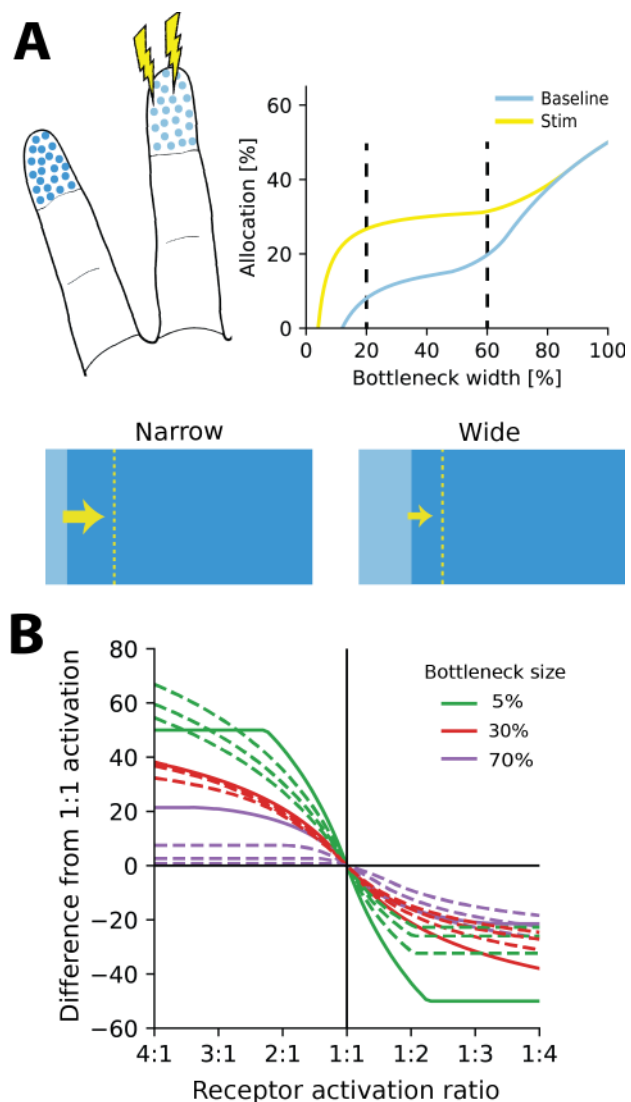


Fig. 4. Re-allocation to account for changes in stimulus statistics. **A.** Top left: Illustration of problem setup. Shown are two fingertips with fixed receptor density and known statistics (baseline condition). Increased stimulation is applied to the middle digit (yellow symbols, stimulation condition), leading to changes in optimal allocations. Top right: Optimal allocations for baseline (blue) and stimulation (yellow) conditions across all bottleneck widths. Stimulation of the middle finger increases its representation, but the relative magnitude of the effect depends on the bottleneck width. Bottom: Changes in allocation of the middle digit for two bottleneck widths (indicated by vertical black dashed lines above). The yellow dotted line denotes increase in allocation post stimulation. This effect is proportionally larger for the narrow compared to the wide bottleneck. **B.** Change in allocation when receptor activation for an input region increases (left half) or decreases (right half), mimicking a change in stimulus statistics. Drastic changes in cortical allocation are observed for narrow bottlenecks (green lines), while wider bottlenecks (red and purple lines) induce more moderate change. Solid lines denote equal receptor density across both regions, while dotted lines show varying density ratios between 1:1 and 1:4.

cused on another predicted outcome of the proposed model: the precise interaction between receptor density and receptor activation in driving resource allocation that we presented earlier. We picked the star-nosed mole as our model system, because stimulus statistics, receptor densities, and cortical allocations have been precisely quantified. Moreover, the star-nosed mole displays considerable variation in all these parameters, presenting a good opportunity to put the model to

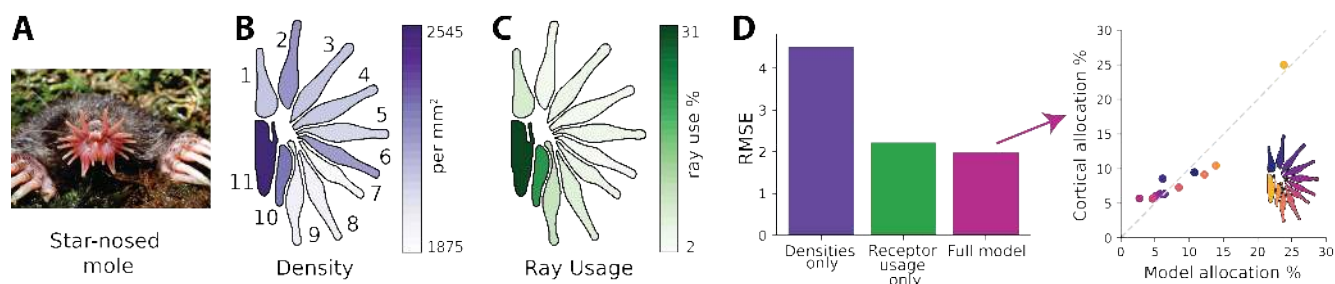


Fig. 5. Resource allocation in the star-nosed mole. **A.** Star-nosed moles have two sets of 11 tactile rays used for detecting and capturing prey. Image by Kenneth Catania and reprinted with permission. **B.** Fiber innervation densities for each ray. **C.** Typical usage percentages for each ray during foraging. Higher usage corresponds to greater average activation of receptors located on the corresponding ray. Typically prey is funneled from distal rays towards ray 11, which is located next to the mouth. Ray outlines reproduced from [Catania and Kaas \(1997\)](#). **D.** Left: root-mean-square error (RMSE) between model predictions and cortical allocation for three different models: restricted to receptor density only (purple), restricted to receptor activation only (green), and a full model (pink) that takes into account both factors. Results confirm previous findings that ray usage is a better predictor of the cortical allocation than receptor densities alone. Additionally, we show that including both of these factors further improves prediction and results in the lowest overall error. Right: predicted versus empirical cortical allocations for all rays. When including both density and activation parameters, the model provides a good fit to empirical measurements.

the test.

The star-nosed mole is a mostly underground dwelling creature relying on active tactile sensing while foraging for prey ([Catania, 2020](#); [Catania and Kaas, 1997](#)). This process is facilitated by two sets of eleven appendages arranged in a star-like pattern that make up the mole's nose (Fig. 5A). Individual rays are tiled with tactile receptors known as Eimer's organs that detect prey, which is then funneled towards the mouth ([Catania and Kaas, 1997](#); [Sawyer and Catania, 2016](#)). The density of fibres innervating the Eimer's organs differs across the rays, with rays closer to the mouth exhibiting higher densities (Figure 5B). The rays also vary in size as well as in their location with respect to the mouth. This affects their usage, as rays closest to the mouth encounter tactile stimuli much more frequently than other rays (Figure 5C). In cortex, a clear topographic representation of the rays can be found ([Catania and Kaas, 1997](#)). However, the extent of the cortical ray representations is not proportional to the physical size of the rays. Ray 11, which sits closest to the mouth, is cortically magnified several-fold, and is considered the tactile equivalent of the visual fovea ([Catania and Remple, 2004](#)).

Previous work by [Catania and Kaas \(1997\)](#) found that cortical sizes are correlated more strongly with the pattern of activation across the rays, rather than their innervation densities. Using the empirical data quantifying receptor densities and stimulus statistics from this study, we investigated whether the efficient coding model could predict typical cortical representation sizes for each ray (see Methods for details), and whether innervation densities or usage would lead to more accurate allocation predictions. As the bottleneck size between periphery and cortex is unknown for the star-nosed mole, we calculated the optimal allocations over all possible bottleneck sizes. Using the model described above we calculated allocations considering density or usage alone, or for both factors combined. We found that the empirical cortical representation sizes are most accurately predicted by a model considering both receptor density and ray usage (Figure 5D,E), suggesting the star-nosed mole could be employing an efficient coding strategy based on decorrelation in the neural represen-

tation of somatosensory inputs.

Discussion

We examined efficient population coding under limited numbers of output neurons in cases of non-uniform input receptor densities and stimulus statistics. Instead of focusing on precise receptive field structure, we asked which coarse region on the sensory sheet the receptive field would fall. We showed that the resulting allocations are nonlinear and depend crucially on the width of the bottleneck, rather than being proportional to the receptor densities or statistics. Specifically, narrow bottlenecks tend to favour expansion of a single region, whereas for larger bottlenecks, allocations converge to a constant ratio between the regions that is closer to a proportional representation. Whether, across all possible bottlenecks, allocations are always expanded, contracted or show both expansion and contraction depends on the relative density and activation ratios, but receptor activation plays a bigger role. When allocation changes due to novel stimulus statistics, a larger fraction of output neurons will switch their receptive field to another region for narrow compared to wide bottlenecks. Finally, we demonstrated that in the star-nosed mole, a model including both accurate innervation densities and statistics of contact provides a better fit to somatosensory cortical region sizes than considering each of these factors alone.

Comparison with previous approaches.

A key feature of efficient coding population models is the non-uniform allocation of output neurons, whereby stimuli occurring at higher probabilities are represented by a greater number of neurons. A common approach is to use Fisher Information as a proxy for mutual information, enabling the calculation of optimal output neuron density and tuning curve placement given a distribution of sensory stimuli ([Ganguli and Simoncelli, 2010, 2014, 2016](#); [Yerxa et al., 2020](#)). However, these approaches maximize information between the output population and the stimulus distribution itself, such that allocating additional neurons to a given region of the input space will always lead to an increase in the overall

amount of information represented. In contrast, our approach assumes a finite number of input receptors in each region. We are thus asking a different question than previous research: once information about a sensory stimulus has been captured in a limited population of receptors, what is the most efficient way of representing this information? This framing implies that once all information from a given region has been fully captured in the output population, our method does not allocate further neurons to that region. Therefore, this also places a limit on the total size of the output population, as this cannot exceed the total number of input receptors.

There has also been prior work on how bottlenecks affect sensory representations, though mostly focused on how different levels of compression affect receptive field structure. For example, [Doi and Lewicki \(2014\)](#) predicted receptive fields of retinal ganglion cells at different eccentricities of the retina, which are subject to different convergence ratios. More recently, such direct effects on representations have also been studied in deep neural networks ([Lindsey et al., 2019](#)). Finally, some approaches employ a different cost function than the mean-squared reconstruction error inherent to the PCA method used here. For example, the Information Bottleneck method ([Tishby et al., 2000](#); [Tishby and Zaslavsky, 2015](#)) aims to find a low-dimensional representation that preserves information about a specific output variable, while compressing information available in the input. It is quite likely that the choice of cost function would affect the resulting allocations, a question that future research should pursue.

Implications for sensory processing.

The quantitative comparison of the model results with the cortical somatosensory representation in the star-nosed mole provided limited evidence that one of the model predictions—namely that receptor density and activation statistics should jointly determine cortical organisation—is borne out in a biological model system. Further direct tests of the theory are hampered by a lack of reliable quantitative empirical data. Nevertheless, the model makes a number of qualitative predictions that can be directly compared to available data or tested in future experiments.

First, there is additional evidence that both receptor density and stimulus statistics drive allocation in neural populations. Similar to the star-nosed model, the primate somatosensory system also exhibits non-uniform distributions of both stimulus statistics and receptor density. Both of these factors are also broadly correlated, for example receptor densities are higher on the fingertips, which are also more likely to be in contact with objects ([Gonzalez et al., 2014](#)). Importantly, while receptor densities alone can explain some of the magnification observed in the cortical somatosensory homunculus, they cannot account for this effect fully ([Corniani and Saal, 2020](#)). Indeed, evidence from non-human primates show that cortical magnification also critically depends on experience ([Xerri et al., 1996, 1999](#)). Similar results have recently been obtained from the brainstem in mice ([Lehnert et al., 2021](#)). Both mechanoreceptor densities ([Verendeev et al., 2015, e.g.](#)) and hand use statistics ([Fragaszy and Crast, 2016](#))

differ across primates, forming the basis for a potential cross-species study. The model presented here demonstrates how both effects can be treated within a single framework driven by information maximization.

Second, it appears that allocation in sensory systems with severely constrained resources is qualitatively in agreement with model predictions. As our results demonstrated, magnification should be more extreme the tighter the bottleneck. The best characterized and most clearly established bottleneck in sensory processing is likely the optical nerve in vision. Given that the optic nerve serves as a narrow bottleneck (approx. 12-27%)¹, and that the fovea contains a much higher density of cone receptors than the periphery ([Curcio et al., 1990](#); [Wells-Gray et al., 2016](#)), the model would predict a large over-representation of the fovea, agreeing with experimental observations (see [Edmondson et al., 2019](#), for further qualitative evidence). In order to further test the proposed model, comparisons could be made between individuals to study variations within a population. Specifically, it would be expected that optic nerves containing a high number of fibers would devote proportionally fewer of them to the fovea than optic nerves containing smaller numbers of fibers (assuming equal receptor densities and numbers in the retina). These comparisons could be extended across species, making use of the fact that photoreceptor densities and optic nerve fiber counts differ across many primate species ([Finlay et al., 2008](#)). Finally, the extent of cortical areas can be controlled by experimental interventions in animals ([Huffman et al., 1999](#)), which would constitute a direct manipulation of the bottleneck.

Re-allocation during development, learning, and ageing.

Changing receptor densities, stimulus statistics, or resource limits over the lifespan of an organism should lead to a dynamic re-allocation of the available resources. The most common case will be changes in stimulus statistics, as both receptor densities and resources should be relatively stable. In such cases, representations should adapt to the new statistics. For example in touch, changing the nature of tactile inputs affects cortical representations ([Coq and Xerri, 1998](#); [Merzenich and Jenkins, 1993](#); [Xerri et al., 1996](#)). Increasing statistics of contact over a region typically leads to expansion of that region in cortex. Our method suggests that the precise level of expansion would be dependent on the bottleneck width with larger effects observed for narrower bottleneck sizes.

For the other cases, changes in fiber numbers during development and ageing might be interpreted as a change in resources. For example, the optic nerve undergoes a period of rapid fiber loss early during development ([Provis et al., 1985](#); [Sefton et al., 1985](#)). Similarly, fiber counts in the optic nerve decrease during ageing ([Dolman et al., 1980](#); [Jonas et al., 1990](#); [Sandell and Peters, 2001](#)). In this case, the model would predict a decrease in the size of peripheral representation in the bottleneck compared to the fovea. It is also pos-

¹Assuming 0.71 - 1.54 million retinal ganglion cells with 80% Midget cells ([Curcio and Allen, 1990](#); [Perry et al., 1984](#)) and 4.6 million cones ([Curcio et al., 1990](#))

sible that the receptor densities themselves may change. In touch, ageing leads to reductions in the densities of receptors in older adults (García-Piqueras et al., 2019). In such cases, we have effectively increased the bottleneck width relative to our receptor population, which again should lead to re-allocation of resources.

Expansion and contraction along the sensory hierarchy.

Magnification of specific sensory input regions can be observed all along the sensory hierarchy, from the brainstem and thalamus up to cortical sensory areas. Often, the amount of expansion and contraction differs between areas. For example, the magnification of the fovea increases along the visual pathway from V1 to V4 (Harvey and Dumoulin, 2011). Which of these representations might be best addressed by the model presented here? We focus on decorrelation, which has been shown to be a driving principle for efficient coding in settings where noise is low (Chalk et al., 2018). This is generally the case in low-level sensory processing and our results might therefore best match early sensory processing up to perhaps low-level cortical representations. Beyond this, it is likely that noise will be high enough for efficient codes to shift away from decorrelation (Hermundstad et al., 2014). Whether and how such a shift would affect neural allocations remains to be seen.

Perceptual consequences. Does the allocation of output neurons lead to testable perceptual consequences? While we do not model neurons' receptive fields directly, allocating more neurons to a given region would increase perceptual spatial acuity for that region. Indeed, cortical magnification and perceptual acuity are correlated in both vision (Duncan and Boynton, 2003) and touch (Duncan and Boynton, 2007). At the same time, the absolute limits on spatial acuity are determined by the density of receptors in each input region. A naive allocation scheme that assigns output neurons proportional to the density of receptors would therefore result in perceptual spatial acuities proportional to receptor distance. Instead, as our results have shown, the allocation should not be proportional in most cases. Specifically, for narrow bottlenecks we would expect relatively higher spatial acuity for regions with high receptor density, than might be expected from a proportional allocation. Conversely, for wider bottlenecks this relationship should be reversed and spatial acuity should be better than expected for lower density regions. In agreement with these results, it has been found that in vision spatial resolution declines faster than expected with increasing eccentricity, suggesting a narrow bottleneck in the optic nerve (Anderson et al., 1991).

A second consequence is that spatial acuity should be better on regions with higher activation probability even when receptor densities are equal. Indeed, spatial discrimination in touch improves with training or even just passive stimulation (Godde et al., 2000; Van Boven et al., 2000), up to a limit that is presumably related to receptor density (Peters et al., 2009; Wong et al., 2013). Assuming a fixed resource limit, training may offer improvements to some digits at the detriment to others. Whether this is indeed the case has to our knowledge

not yet been empirically tested.

Finally, previous work has shown that non-uniform tuning curves across a population will lead to characteristic biases in perceptual tasks (Wei and Stocker, 2015). While the original formulation assumed that this heterogeneous allocation of output neurons was driven by stimulus statistics alone, we have shown here that it can also be a consequence of receptor densities. Thus, perceptual biases might also be expected to arise from neural populations that efficiently represent sensory inputs sampled by non-uniform receptor populations.

Limitations and future work.

We considered noiseless, second-order, linear models only and demonstrated that even such simple models exhibit surprising complexity in how they manage trade-offs in resource allocation under constraints. Specifically, we found that output neurons were not generally allocated proportionally to input neurons or according to some other fixed rule. It therefore stands to reason that similarly complex trade-offs would manifest in more complex models, even though the precise allocations might differ. Nevertheless, since principal component analysis is widely employed for decorrelation and dimensionality reduction, and therefore incorporated into many other algorithms, our results immediately generalize to several other methods. This includes independent component analysis, which considers higher-order rather than second-order statistics, but relies on a whitened signal, which in the undercomplete (bottleneck) case is obtained via PCA (Hyvärinen and Oja, 2000). Similarly, some models that do incorporate sensory noise and maximize reconstruction accuracy also use an undercomplete set of principal components to reduce the dimensionality of the sensory signal (Doi and Lewicki, 2014). In both of these examples the resulting receptive field structure will differ, but their allocation—where on the sensory sheet they will fall—will be governed by the same principles described earlier.

We derived allocations for simple negative exponential covariance functions. Choosing this parametric function allowed us to calculate allocation and limits analytically for large populations under certain circumstances. While it is possible to solve the allocation problem numerically for arbitrary covariance functions, we noticed that the presence of small numerical errors can affect the sorting process and caution is therefore warranted. Negative exponential covariance functions obey the intuitive notion that receptor correlations decrease with distance, and they are likely a good match for the sense of touch. However, alternative covariance functions do occur in sensory systems, with the most prominent example given by the visual system, where eigenvalue decay follows a well-established power law (Field, 1987). In such cases, precise allocation will differ, however as long as the covariance function decreases with distance, our results should hold qualitatively (see Edmondson et al., 2019, for numerical calculations using a covariance functions empirically measured from natural images).

Methods

Our main goal was to derive a method for efficiently allocating output neurons to one of several input regions with different correlational response structure, in the presence of constraints on the number of output neurons or amount of information being transmitted. In the following, we focus on the main rationale and equations, while specific proofs can be found in the appendix. First, we outline the framework for combined whitening and dimensionality reduction that is employed. Next, we demonstrate how this framework can be applied to multiple input regions with different statistics and densities of receptors, and how calculation of the eigenvalues of region-specific covariance matrices solves the problem of resource allocation. Finally, we demonstrate how the problem can be solved analytically for a certain choice of covariance function.

Combined whitening and dimensionality reduction.

We assume that receptors are arranged on a two-dimensional sensory sheet. Correlations in the inputs occur as receptors that are nearby in space have more similar responses. Information is maximised in such a setup by decorrelating the sensory inputs. Here we decorrelate using a second-order linear model. To model the bottleneck, we restrict the outputs to $m < n$, where n is the total number of receptors.

If the inputs are represented as a matrix \mathbf{X} of dimensions $n \times z$ (where z is the number of sensory input patterns), then our goal is to find an $m \times n$ dimensional matrix \mathbf{W} such that $\mathbf{W}\mathbf{X}$ is uncorrelated:

$$\mathbf{X}^T \mathbf{W}^T \mathbf{W} \mathbf{X} = \mathbf{I}. \quad (1)$$

This is achieved by setting $\mathbf{W} = \mathbf{\Sigma}^{-\frac{1}{2}}$, where $\mathbf{\Sigma} = \mathbf{X}^T \mathbf{X}$. Solutions can then be expressed in terms of the diagonal matrix of eigenvalues, $\mathbf{\Lambda}$, and eigenvectors, $\mathbf{\Phi}$, of the covariance matrix $\mathbf{\Sigma}$:

$$\mathbf{W} = \mathbf{P} \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{\Phi}^T. \quad (2)$$

Whitening filters are not unique and obtained with any orthogonal matrix \mathbf{P} . If $\mathbf{P} = \mathbf{I}$, the receptive fields \mathbf{W} are the principal components of the inputs (standard PCA). Localized receptive fields are obtained by setting $\mathbf{P} = \mathbf{\Phi}$, which is known as ZCA (Zero-Phase Component Analysis). In cases with a bottleneck the solution involves solving an Orthogonal Procrustes problem (Doi and Lewicki, 2014) to find \mathbf{P}^* , an m -dimensional orthogonal matrix (where m is the size of the bottleneck) which minimizes the reconstruction error of the inputs and a set of ideal local receptive fields \mathbf{W}_{opt} :

$$\mathbf{P}^* = \min_{\mathbf{P}} \left\| \mathbf{W}_{opt} - \mathbf{P} \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{\Phi}^T \right\|_F^2, \quad (3)$$

where $\|\cdot\|_F$ denotes the Frobenius norm, and $\mathbf{\Lambda}$ and $\mathbf{\Phi}$ are as above but retaining only those components with the m largest eigenvalues. In either case it can be seen that independent of the precise receptive field structure (which will partly depend on \mathbf{P}), the solution relies on the m largest eigenvalues of the covariance matrix $\mathbf{\Sigma}$. It is this aspect that we will exploit in

the following and that allows us to determine the input region within which a receptive field falls.

Extension to multiple input regions.

For our specific problem, we are interested in the case of multiple input regions with different correlational structure (i.e. due to differing receptor density or activation). To simplify the derivations we approximate different input regions as independent, such that the overall covariance matrix will be a block diagonal matrix. The covariance $\mathbf{\Sigma}$ for two input regions, R_1 and R_2 , can then be expressed as follows:

$$\mathbf{\Sigma} = \begin{bmatrix} \mathbf{\Sigma}^{(R1)} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Sigma}^{(R2)} \end{bmatrix} \quad (4)$$

This assumption turns out to be a reasonable approximation when region sizes are relatively big and correlations typically do not extend far across the sensory sheet (see Edmondson et al., 2019, for a comparison between block and non-block region covariance matrices in 1D). Furthermore, in many sensory systems the borders between regions of differing density tend to be relatively narrow. For example, in touch, the digits of the hand are spatially separated, and regions of differing densities, for example between the digit tips and proximal phalanges, neighbour along the short rather than long axis of the digit. In the star-nosed mole, rays of different innervation densities are separated and neighbour only along their connection to the rest of nose. However, the block matrix approximation might be problematic in cases with many very small adjacent regions with strong, far-ranging correlations. The eigenvalues and eigenvectors of a block diagonal covariance matrix also follow the block diagonal form and can be calculated from the individual region covariances alone, by a simple application of the Cauchy interlacing theorem. Thus, the corresponding eigenvalues and eigenvectors are:

$$\mathbf{\Lambda} = \begin{bmatrix} \mathbf{\Lambda}^{(R1)} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Lambda}^{(R2)} \end{bmatrix} \quad \text{and} \quad \mathbf{\Phi} = \begin{bmatrix} \mathbf{\Phi}^{(R1)} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Phi}^{(R2)} \end{bmatrix} \quad (5)$$

Due to the sensory bottleneck, only the m largest eigenvalues from the combined set $\mathbf{\Lambda}$ will be retained. An eigenvalue selected from $\mathbf{\Lambda}^{(R1)}$ indicates that the receptive field of the corresponding output neuron will fall onto region R_1 .

Calculation of eigenvalues.

We model the covariance between receptors as a negative exponential function. The covariance matrix is then calculated as a function of distance between pairs of receptors as follows (see Fig. 1C). For region R_1 we get:

$$\Sigma_{ij}^{(R1)} = e^{-\sigma|x_i-x_j|}, \quad (6)$$

where x_i and x_j are the locations of the i th and j th receptors, and σ is the decay constant. To calculate the covariance for R_2 , we need to take into account the potentially different receptor density and response variance for this region. Denoting the ratio of the response variances between both regions by a , and the ratio of receptor densities by d , the covariance for R_2 can be expressed as:

$$\Sigma_{ij}^{(R2)} = ae^{-d\sigma|x_i-x_j|}. \quad (7)$$

It can be seen that a scales the overall covariance matrix (see also Fig. 1C), while d changes the spatial extent of the correlations and thereby implicitly accounts for the different receptor density.

The corresponding eigenvalues for an exponential covariance function in the continuous domain can be calculated analytically. The eigenvalue-eigenvector problem is expressed as an integral homogeneous equation, such that for R_1 we get:

$$\lambda_k \phi_k(x) = \int_0^L e^{-\sigma|x-y|} \phi_k(y) dy, \quad (8)$$

where $\phi_k(x)$ is the k th eigenfunction and λ_k its corresponding eigenvalue. The domain length L is the input region size for one of the dimensions.

It can be shown that solutions to this problem can be related to the Laplacian operator (see appendices B and C for proofs), such that:

$$\lambda_k = \frac{2\sigma}{\mu_k + \mu_k + \sigma^2}, \quad (9)$$

where μ_k are the eigenvalues of the Laplacian operator. The general solution for the Laplacian eigenvalue problem for a two-dimensional rectangle with Dirichlet boundary conditions is (Strauss, 2007):

$$\mu_k = \mu_{l,m} = \frac{l^2 \pi^2}{L_1^2} + \frac{m^2 \pi^2}{L_2^2}, l, m = 1, 2, \dots \quad (10)$$

where L_1 and L_2 are the size of the domain for each dimension. The calculation of the eigenvalues for R_2 proceeds analogously. The eigenvalue-eigenvector problem is given as:

$$\lambda_k \phi_k(x) = \int_0^{L\sqrt{d}} ae^{-d\sigma|x-y|} \phi_k(y) dy \quad (11)$$

The total density of a square region, d , is calculated as $(\sqrt{d})^2 = d$, where each axis of the region has density \sqrt{d} . Since the receptor density ratio d causes an implicit stretching of the space, the region length L needs to be adjusted accordingly. In 2D, each axis is therefore scaled by \sqrt{d} , so that $L = L\sqrt{d}$.

Allocation in the bottleneck.

Given a sensory system with limited representational capacity, different regions may be allocated different amounts of resources. Here we calculate the allocations over different bottleneck widths for two regions, while the extension to multiple regions is given in appendix E. In the following, we assume two-dimensional square regions of equal size for ease of analysis (see appendix F for the equivalent solution in 1D). A single variable L is therefore used to denote the lengths of the squares. Following (9), the eigenvalues for regions R_1 and R_2 can now be calculated as:

$$R_1: \lambda_{l,m}^{(R1)} = \frac{2\sigma}{\pi^2 L^{-2}(l^2 + m^2) + \sigma^2} \quad (12)$$

$$R_2: \lambda_{n,o}^{(R2)} = \frac{2\sigma a \sqrt{d}}{\pi^2 L^{-2}(n^2 + o^2) + \sigma^2} \quad (13)$$

where l, m and $n, o \in \mathbb{N}$ enumerate different eigenvalues for regions R_1 and R_2 , respectively.

In order to calculate how many output neurons are allocated to R_1 and R_2 for different bottleneck widths, we will need to establish an ordering of the eigenvalues, such that for each pair (l, m) we can determine the sorted rank of the eigenvalues. In contrast to the 1D case (see F), there is no natural ordering of the eigenvalues in two dimensions, however a close approximation can be obtained by calculating the number of lattice points enclosed by a quarter circle with radius $p = l^2 + m^2$ (see appendix D for full details). Denoting this function as $N(p)$ and setting $p^{(R1)} = l^2 + m^2$ and $p^{(R2)} = n^2 + o^2$, we can then calculate the number of eigenvalues allocated to R_1 as a function of the number of neurons allocated to R_2 , by setting $\lambda^{(R1)} = \lambda^{(R2)}$ and solving for $p^{(R2)}$. This yields:

$$p^{(R2)} = a\sqrt{d}p^{(R1)} + \frac{L^2\sigma^2 a\sqrt{d} - L^2\sigma^2}{\pi^2}. \quad (14)$$

As we allocate more neurons to region R_1 , the ratio $\frac{N(p^{(R1)})}{N(p^{(R2)})}$ simplifies to: $\lim_{R_1 \rightarrow \infty} \frac{N(p^{(R1)})}{N(p^{(R2)})} = a\sqrt{d}$. The fraction of neurons allocated to each region therefore depends on the size of the bottleneck and converges to $\frac{1}{1+a\sqrt{d}}$ and $\frac{a\sqrt{d}}{1+a\sqrt{d}}$ for R_1 and R_2 respectively.

Calculations for star-nosed mole.

The eleven rays were approximated as 2D square regions with areas set to their reported sizes (Sawyer and Catania, 2016). Receptor densities for each ray were calculated as the sensory fiber innervation per mm^2 (Catania and Kaas, 1997). Approximations of receptor activation on each ray were calculated from empirical data of prey foraging interactions recorded by Catania and Kaas (1997). Contact statistics were converted to receptor activation probabilities with receptors following a Bernoulli distribution. Finally, activation variance was calculated as the variance of the Bernoulli distribution (see also appendix A). The decay rate σ of the negative exponential covariance function was determined for each ray using a model of the typical extent of receptor activation during interaction with prey stimuli of varying size. Each ray interacts with varying prey sizes at different frequencies. For example, ray 11 is typically contacted by smaller stimuli more often than other rays. A 2D model of the rays was used to simulate average responses to each stimulus size. Each model ray was tiled with receptors, and circular stimuli of different sizes were then randomly placed over the ray. The radii and frequencies of each stimulus size were based on the prey model (Catania and Kaas, 1997)). A ray receptor was marked as active if its coordinate position was within the bounds of the stimuli. Response covariance between receptors was then calculated and an exponential function was fit to find the σ decay parameter. To determine allocations,

the eigenvalues for each ray were calculated analytically using equation (A.24), resulting in allocations for each ray at all bottleneck widths. Three models were compared: first, a 'density only' model, which includes accurate receptor densities, but receptor activation remains uniform across all rays; second, an 'activation only' model, which includes heterogeneous receptor activation, but uniform receptor density across all rays; finally, the 'full model' combines both accurate densities and receptor activation. Model allocations for each ray were compared to the cortical allocation empirical data from Catania and Kaas (1997). As the bottleneck size for the star-nosed mole is unknown, the root-mean-square error (RMSE) was calculated for each model at all bottleneck widths. The bottleneck resulting in the lowest error was then selected for each. For the activation only and full models, the lowest RMSE values were between 37–45% bottleneck widths. For the density only model, RMSE was consistently between 4.5–5 over all bottleneck widths.

ACKNOWLEDGEMENTS

This work was supported by the Wellcome Trust [209998/Z/17/Z] and by the European Union Horizon 2020 program as part of the Human Brain Project [HBP-SGA2, 785907]. The manuscript was typeset using the Henriques lab bioRxiv template.

Bibliography

- Anderson SJ, Mullen KT, Hess RF. Human peripheral spatial resolution for achromatic and chromatic stimuli: limits imposed by optical and retinal factors. *J. Physiol.* 442: 47–64, 1991.
- Atick JJ. Could information theory provide an ecological theory of sensory processing? *Network: Computation in neural systems* 3: 213–251, 1992.
- Atick JJ, Redlich AN. Towards a theory of early visual processing. *Neural Computation* 2: 308–320, 1990.
- Attneave F. Some informational aspects of visual perception. *Psychological review* 61: 183, 1954.
- Azzopardi P, Cowey A. Preferential representation of the fovea in the primary visual cortex. *Nature* 361: 719–721, 1993.
- Barlow HB. Possible principles underlying the transformation of sensory messages. *Sensory communication* 1: 217–234, 1961.
- Bell AJ, Sejnowski TJ. The “independent components” of natural scenes are edge filters. *Vision Research* 37: 3327–3338, 1997.
- Catani M. A little man of some importance. *Brain* 140: 3055–3061, 2017.
- Catania KC. All in the family—touch versus olfaction in moles. *The Anatomical Record* 303: 65–76, 2020.
- Catania KC, Kaas JH. Somatosensory fovea in the star-nosed mole. *Journal of Comparative Neurology* 387: 215–233, 1997.
- Catania KC, Remple FE. Tactile foveation in the star-nosed mole. *Brain, behavior and evolution* 63: 1–12, 2004.
- Chalk M, Marre O, Tkačik G. Toward a unified theory of efficient, predictive, and sparse coding. *Proceedings of the National Academy of Sciences* 115: 186–191, 2018.
- Chechik G, Anderson MJ, Bar-Yosef O, Young ED, Tishby N, Nelken I. Reduction of information redundancy in the ascending auditory pathway. *Neuron* 51: 359–368, 2006.
- Coq JO, Xerri C. Environmental enrichment alters organizational features of the forepaw representation in the primary somatosensory cortex of adult rats. *Experimental brain research* 121: 191–204, 1998.
- Corniani G, Saal HP. Tactile innervation densities across the whole body. *J. Neurophysiol.* 124: 1229–1240, 2020.
- Curcio CA, Allen KA. Topography of ganglion cells in human retina. *Journal of comparative Neurology* 300: 5–25, 1990.
- Curcio CA, Sloan KR, Kalina RE, Hendrickson AE. Human photoreceptor topography. *Journal of comparative neurology* 292: 497–523, 1990.
- Doi E, Gauthier JL, Field GD, Shlens J, Sher A, Greschner M, Machado TA, Jepsen LH, Mathieson K, Gunning DE et al. Efficient coding of spatial information in the primate retina. *Journal of Neuroscience* 32: 16256–16264, 2012.
- Doi E, Lewicki MS. A simple model of optimal population coding for sensory systems. *PLoS computational biology* 10: e1003761, 2014.
- Dolman CL, McCormick AQ, Drance SM. Aging of the optic nerve. *Archives of ophthalmology* 98: 2053–2058, 1980.
- Duncan RO, Boynton GM. Cortical magnification within human primary visual cortex correlates with acuity thresholds. *Neuron* 38: 659–671, 2003.
- Duncan RO, Boynton GM. Tactile hyperacuity thresholds correlate with finger maps in primary somatosensory cortex (s1). *Cerebral Cortex* 17: 2878–2891, 2007.
- Edmondson LR, Rodriguez AJ, Saal HP 2019 Nonlinear scaling of resource allocation in sensory bottlenecks In: *Advances in Neural Information Processing Systems*, p. 7545–7554.
- Engel SA, Glover GH, Wandell BA. Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cerebral cortex (New York, NY: 1991)* 7: 181–192, 1997.
- Field DJ. Relations between the statistics of natural images and the response properties of cortical cells. *Josa a 4*: 2379–2394, 1987.
- Finlay BL, Franco ECS, Yamada ES, Crowley JC, Parsons M, Muniz JAP, Silveira LCL. Number and topography of cones, rods and optic nerve axons in new and old world primates. *Visual Neuroscience* 25: 289, 2008.
- Fragaszy DM, Crast J. Functions of the hand in primates In: *The evolution of the primate hand*. Springer, 2016 p. 313–344.
- Ganguli D, Simoncelli E. Implicit encoding of prior probabilities in optimal neural populations. *Advances in neural information processing systems* 23: 658–666, 2010.
- Ganguli D, Simoncelli EP. Efficient sensory encoding and bayesian inference with heterogeneous neural populations. *Neural computation* 26: 2103–2134, 2014.
- Ganguli D, Simoncelli EP. Neural and perceptual signatures of efficient sensory coding. *arXiv preprint arXiv:1603.00058*, 2016.
- García-Piqueras J, García-Mesa Y, Cárcaba L, Feito J, Torres-Parejo I, Martín-Biedma B, Cobo J, García-Suárez O, Vega JA. Ageing of the somatosensory system at the periphery: age-related changes in cutaneous mechanoreceptors. *Journal of anatomy* 234: 839–852, 2019.
- Godde B, Stauffenberg B, Spengler F, Dinse HR. Tactile coactivation-induced changes in spatial discrimination performance. *Journal of Neuroscience* 20: 1597–1604, 2000.
- Gonzalez F, Gosselin F, Bachtá W. Analysis of hand contact areas and interaction capabilities during manipulation and exploration. *IEEE transactions on haptics* 7: 415–429, 2014.
- Goodchild AK, Ghosh KK, Martín PR. Comparison of photoreceptor spatial density and ganglion cell morphology in the retina of human, macaque monkey, cat, and the marmoset callithrix jacchus. *Journal of Comparative Neurology* 366: 55–75, 1996.
- Graham DJ, Chandler DM, Field DJ. Can the theory of “whitening” explain the center-surround properties of retinal ganglion cell receptive fields? *Vision Research*. 46: 2901–2913, 2006.
- Graham D, Field D. Natural images: Coding efficiency. *Encyclopedia of Neuroscience* 6: 19–27, 2009.
- Harvey BM, Dumoulin SO. The relationship between cortical magnification factor and population receptive field size in human visual cortex: Constancies in cortical architecture. *Journal of Neuroscience* 31: 13604–13612, 2011.
- Hermundstadt AM, Briguglio JJ, Conte MM, Victor JD, Balasubramanian V, Tkacik G. Variance predicts saliency in central sensory processing. *Life*, 2014.
- Huffman KJ, Molnar Z, Van Dellen A, Kahn DM, Blakemore C, Krubitzer L. Formation of cortical fields on a reduced cortical sheet. *J. Neurosci.* 19: 9939–9952, 1999.
- Hyvärinen A, Oja E. Independent component analysis: algorithms and applications. *Neural networks* 13: 411–430, 2000.
- Jenkins WM, Merzenich MM, Ochs MT, Allard T, Guic-Robles E. Functional reorganization of primary somatosensory cortex in adult owl monkeys after behaviorally controlled tactile stimulation. *Journal of neurophysiology*, 1990.
- Johansson RS, Vallbo A. Tactile sensibility in the human hand: relative and absolute densities of four types of mechanoreceptive units in glabrous skin. *The Journal of physiology* 286: 283–300, 1979.
- Jonas JB, Müller-Bergh J, Schlötzer-Schrehardt U, Naumann G. Histomorphometry of the human optic nerve. *Investigative ophthalmology & visual science* 31: 736–744, 1990.
- Karklin Y, Simoncelli EP. Efficient coding of natural images with a population of noisy linear-nonlinear neurons. *Advances in neural information processing systems* 24: 999, 2011.
- Kersten D. Predictability and redundancy of natural images. *J. Opt. Soc. Am. A*: 2395, 1987.
- Lehert BP, Santiago C, Huey EL, Emanuel AJ, Renaud S, Africawala N, Alkisar I, Zheng Y, Bai L, Koutsoumpa C, Hong JT, Magee AR, Harvey CD, Ginty DD 2021 Mechanoreceptor synapses in the brainstem shape the central representation of touch unpublished.
- Lewicki MS. Efficient coding of natural sounds. *Nature neuroscience* 5: 356–363, 2002.
- Lindsey J, Ocko SA, Ganguli S, Deny S. A unified theory of early visual representations from retina to cortex through anatomically constrained deep CNNs. *International Conference on Learning Representations (ICLR)*, 2019.
- Martuzzi R, van der Zwaag W, Farthouat J, Gruetter R, Blanke O. Human finger somatotopy in areas 3b, 1, and 2: a 7t fmri study using a natural stimulus. *Human brain mapping* 35: 213–226, 2014.
- Merzenich MM, Jenkins WM. Reorganization of cortical representations of the hand following alterations of skin inputs induced by nerve injury, skin island transfers, and experience. *Journal of hand therapy* 6: 89–104, 1993.
- Olshausen BA, Field DJ. Wavelet-like receptive fields emerge from a network that learns sparse codes for natural images. *Nature* 381: 607–609, 1996.
- Olshausen BA, Field DJ. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision research* 37: 3311–3325, 1997.
- Olshausen BA, Field DJ. Sparse coding of sensory inputs. *Current opinion in neurobiology* 14: 481–487, 2004.
- Perry V, Oehler R, Cowey A. Retinal ganglion cells that project to the dorsal lateral geniculate nucleus in the macaque monkey. *Neuroscience* 12: 1101–1123, 1984.
- Peters RM, Hackeman E, Goldreich D. Diminutive digits discern delicate details: fingertip size and the sex difference in tactile spatial acuity. *Journal of Neuroscience* 29: 15756–15761, 2009.
- Provis JM, Van Driel D, Bilson FA, Russell P. Human fetal optic nerve: overproduction and elimination of retinal axons during development. *Journal of Comparative Neurology* 238: 92–100, 1985.
- Ritter H, Martinez T, Schulten K. *Neural Computation and Self-organizing Maps: An Introduction* Addison-Wesley, 1992.
- Ritter H, Schulten K. On the stationary state of kohonen’s self-organizing sensory mapping. *Biological cybernetics* 54: 99–106, 1986.
- Sandell JH, Peters A. Effects of age on nerve fibers in the rhesus monkey optic nerve. *Journal of Comparative Neurology* 429: 541–553, 2001.
- Sawyer EK, Catania KC. Somatosensory organ topography across the star of the star-nosed mole (condylura cristata). *Journal of Comparative Neurology* 524: 917–929, 2016.
- Sefton AJ, Horsburgh GM, Lam K. The development of the optic nerve in rodents. *Australian and New Zealand journal of ophthalmology* 13: 135–145, 1985.
- Sereno MI, Dale A, Reppas J, Kwong K, Belliveau J, Brady T, Rosen B, Tootell R. Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Sci-*

ence 268: 889–893, 1995.

Shwartz-Ziv R, Tishby N. Opening the black box of deep neural networks via information. *arXiv preprint arXiv:1703.00810*, 2017.

Smith EC, Lewicki MS. Efficient auditory coding. *Nature* 439: 978–982, 2006.

Strauss WA. *Partial differential equations: An introduction* John Wiley & Sons, 2007.

Tishby N, Pereira FC, Bialek W. The information bottleneck method. *arXiv preprint physics/0004057*, 2000.

Tishby N, Zaslavsky N 2015 Deep learning and the information bottleneck principle In: *2015 IEEE Information Theory Workshop (ITW)*, p. 1–5. IEEE.

Van Boven RW, Hamilton RH, Kauffman T, Keenan JP, Pascual-Leone A. Tactile spatial resolution in blind braille readers. *Neurology* 54: 2230–2236, 2000.

Verendeef A, Thomas C, McFarlin SC, Hopkins WD, Phillips KA, Sherwood CC. Comparative analysis of meissner's corpuscles in the fingertips of primates. *Journal of anatomy* 227: 72–80, 2015.

Wei XX, Stocker AA. A bayesian observer model constrained by efficient coding can explain "anti-bayesian" percepts. *Nature Neuroscience* 18: 1509, 2015.

Wells-Gray E, Choi S, Bries A, Doble N. Variation in rod and cone density from the fovea to the mid-periphery in healthy human retinas using adaptive optics scanning laser ophthalmoscopy. *Eye* 30: 1135, 2016.

Wong M, Peters RM, Goldreich D. A physical constraint on perceptual learning: Tactile spatial acuity improves with training to a limit set by finger size. *Journal of Neuroscience* 33: 9345–9352, 2013.

Xerri C, Coq J, Merzenich M, Jenkins W. Experience-induced plasticity of cutaneous maps in the primary somatosensory cortex of adult monkeys and rats. *Journal of Physiology-Paris* 90: 277–287, 1996.

Xerri C, Merzenich MM, Jenkins W, Santucci S. Representational plasticity in cortical area 3b paralleling tactual-motor skill acquisition in adult monkeys. *Cerebral cortex* 9: 264–276, 1999.

Yarrow S, Challis E, Seriès P. Fisher and shannon information in finite neural populations. *Neural computation* 24: 1740–1780, 2012.

Yerxa TE, Kee E, DeWeese MR, Cooper EA. Efficient sensory coding of multidimensional stimuli. *PLoS computational biology* 16: e1008146, 2020.

Appendix 1: Mathematical proofs and derivations

A. Stimulus statistics and response variance.

Decorrelation works on second-order statistics and therefore stimulus statistics would only be taken into account by the model if they affect the covariance matrix. One way this can happen is through the extent of the spatial correlations (parameter σ in the covariance function). For example, in touch the size distribution of stimuli that would typically make contact with a given skin region might differ, leading to a different correlational structure. While we calculate allocations for different values of σ , we keep this value fixed across different input regions, for simplicity.

A more common case is that of receptors in one region being more active than receptors in another region. For example, in touch the fingertips make contact with objects much more frequently than does the palm. An increased probability of making contact would translate into higher receptor response rates. In turn, higher response rates imply higher response variance, which would be directly reflected in the covariance matrix as a multiplicative scaling of the covariance function. For example, assuming that each receptor follows a simple Bernoulli distribution (either on or off, with a probability of p being on), then the response variance can be calculated as $p(1-p) = p - p^2$. Assuming that the likelihood of any receptor being active is generally low, the variance scales almost linearly with receptor activation. Differences in activation between two regions are represented by the activation ratio a throughout the paper.

B. Relationship between PCA and Laplacian eigenvalue problem.

B.1. Rationale.

Let Ω be a region with a density of receptors ρ . In a one dimensional region $\Omega = [0, L]$, the density can be expressed as $\rho = \frac{N}{L}$ or number of receptors per unit length. Assuming an exponential decay of correlations, the covariance between receptors i and j is:

$$C(i, j) = e^{(-\sigma|i\Delta x - j\Delta x|)}, \quad (\text{A.1})$$

where $\Delta x = 1/\rho$ is the distance between receptors. Subsampling the space by taking a fraction N/d of the original receptors, $d > 1$, the covariance for positions i, j , becomes:

$$\tilde{C}(i, j) = e^{(-\sigma|id\Delta x - jd\Delta x|)} = e^{(-d\sigma|i\Delta x - j\Delta x|)}. \quad (\text{A.2})$$

Therefore, we encode the changes in receptor density in a scaling of the exponential decay rate. For a given distribution of receptors, there is an induced partition of the interval $[0, L]$, therefore, for a fixed $x = i\Delta x$, the covariance in the j th bin is approximately equal to the area of the exponential covered in that bin:

$$\int_{(j-1)\Delta x}^{j\Delta x} C(x, y) dy \approx C(i\Delta x, j\Delta x) \Delta x,$$

Summing over all the bins, we arrive at the PCA problem:

$$\sum_{j=0}^n C(i\Delta x, j\Delta x) \phi(j\Delta x) \Delta x = \lambda \phi(i\Delta x). \quad (\text{A.3})$$

The continuum limit is found formally as $\Delta x \rightarrow 0$.

B.2. Derivation.

In order to find the optimal assignment for a given receptor density, we are interested in solutions to the following equation which can be seen as a continuous version of the traditional PCA problem with an exponentially decaying covariance matrix:

$$\lambda \phi(x) = \int_0^L e^{(-\sigma|x-y|)} \phi(y) dy, \quad (\text{A.4})$$

where σ is the decay rate. We are interested in solutions $\phi \in C^2(\mathbb{R})$, that is, twice differentiable solutions that satisfy appropriate boundary conditions.

Theorem 1: If ϕ is a solution of (A.4), then it is an eigenfunction of the Laplacian operator with eigenvalues:

$$\mu = \frac{2\sigma}{\lambda} - \sigma^2. \quad (\text{A.5})$$

That is, in one dimension, solutions ϕ satisfy:

$$-\frac{d^2}{dx^2} \phi(x) = \mu \phi(x). \quad (\text{A.6})$$

Proof: Differentiating equation (A.4) twice using the Leibniz rule we obtain:

$$\frac{d}{dx} \phi(x) = \frac{\sigma}{\lambda} \left\{ -\int_0^x e^{-\sigma(x-y)} \phi(y) dy + \int_x^L e^{\sigma(x-y)} \phi(y) dy \right\}, \quad (\text{A.7})$$

$$\frac{d^2}{dx^2} \phi(x) = \frac{\sigma}{\lambda} \left\{ -2\phi(x) + \sigma \int_0^L e^{-\sigma|x-y|} \phi(y) dy \right\}.$$

The second term on the right hand side can be replaced using (A.4) obtaining the desired result:

$$\frac{d^2}{dx^2} \phi(x) = -\frac{2\sigma}{\lambda} \phi(x) + \sigma^2 \phi(x)$$

or

$$-\frac{d^2}{dx^2} \phi(x) = \left(\frac{2\sigma}{\lambda} - \sigma^2 \right) \phi(x). \quad \blacksquare$$

The previous is a sufficient conditions on the solutions to (A.4). A necessary condition is given in the following theorem:

Theorem 2: A solution to (A.6) is a also solution to (A.4) if it satisfies the following boundary conditions:

$$\phi'(0) = \sigma \phi(0) \quad (\text{A.8})$$

$$\phi'(L) = -\sigma \phi(L). \quad (\text{A.9})$$

Proof: Let's assume ϕ is a solution to (A.6). We proceed by convolving (A.6) on both sides with the kernel $e^{-\sigma x}$:

$$\int_0^x e^{-\sigma(x-y)} \phi''(y) dy = \mu \int_0^x e^{-\sigma(x-y)} \phi(y) dy.$$

Integrating by parts twice we get:

$$\begin{aligned} \phi'(x) - e^{-\sigma x} \phi'(0) - \sigma \phi(x) + \sigma e^{-\sigma x} \phi(0) + \\ \sigma^2 \int_0^x e^{-\sigma(x-y)} \phi(y) dy \\ = \mu \int_0^x e^{-\sigma(x-y)} \phi(y) dy. \end{aligned} \quad (\text{A.10})$$

Using (A.5) and (A.9), we obtain:

$$-\phi'(x) + \sigma \phi(x) = \frac{2\sigma}{\lambda} \int_0^x e^{-\sigma(x-y)} \phi(y) dy. \quad (\text{A.11})$$

Repeating the procedure with the kernel $e^{\sigma x}$ in the interval $[x, L]$, yields:

$$\phi'(x) + \sigma \phi(x) = \frac{2\sigma}{\lambda} \int_x^L e^{-\sigma(x-y)} \phi(y) dy. \quad (\text{A.12})$$

Adding (A.11) and (A.12) we recover (A.4) which finalizes the proof. ■

B.3. Scaling.

In some instances of our problem, the exponential covariance will be scaled by the activation ratio, a . In general, the same reasoning applies to any linear combination of solutions, therefore, our results extend to that case. In particular, we have the following result:

Theorem 3: The eigenvalues of the scaled covariance matrix, $C'(x, y) = aC(x, y)$ are:

$$\lambda_s = a\lambda, \quad (\text{A.13})$$

where λ is an eigenvalue of the original problem.

Proof: Let $\phi(x)$ be a solution of (A.4). By linearity of the integral we have:

$$\begin{aligned} \int_0^L a e^{(-\sigma|x-y|)} \phi(y) dy \\ = a \int_0^L e^{(-\sigma|x-y|)} \phi(y) dy = a\lambda \phi(x) \end{aligned} \quad (\text{A.14})$$

■

C. Solutions.

In the previous section we saw that solutions of the PCA problem (A.4) and Laplacian (A.6) coincide if boundary conditions (A.9) and (A.8) are met.

In this section we show how these solutions relate to solutions of the boundary value problem of (A.6) with $\phi(0) = \phi(L) = 0$, which correspond to the eigenmodes of an idealized vibrating string fixed at the extremes. Such modes are considerably simpler than the exact ones and, as we show, are sufficient for our analysis.

C.1. Eigenmodes of a vibrating string.

For $\kappa = \sqrt{\mu}$, $\mu > 0$, solutions can be found by assuming a general solution of the form:

$$\phi(x) = B \sin(\kappa x) + C \cos(\kappa x), \quad (\text{A.15})$$

and then using the boundary conditions to find the constants. The first boundary condition implies that $C = 0$. The second boundary condition gives the equation:

$$\sin(\kappa L) = 0$$

which is satisfied for $\kappa L = n\pi$ or

$$\mu = \frac{n^2 \pi^2}{L^2}. \quad (\text{A.16})$$

where $n = 1, 2, \dots$ is the index of eigenvalue.

C.2. Exact eigenvalues.

In order to find the exact analytical eigenvalues of (A.4), we again assume $\mu > 0$ and a solution of the form (A.15). Using the first boundary condition (A.8, A.9) we get the following relationship:

$$B = \frac{\sigma}{\kappa} C, \quad (\text{A.17})$$

and with the second boundary condition, we obtain

$$\tan \kappa L = \frac{\kappa B + \sigma C}{\kappa C - \sigma B}; \quad (\text{A.18})$$

replacing (A.17), we find the transcendental equation

$$\tan \kappa L = \frac{2\sigma \kappa}{\kappa^2 - \sigma^2}, \quad (\text{A.19})$$

whose solutions lead to the exact eigenvalues of the continuous PCA problem. Replacing (A.16) in (A.19) and reorganizing the right hand side, we find that the regimes for which our approximation is most accurate, are those for which

$$\left(\frac{2n\pi}{L\sigma} \right) \left(\frac{1}{\left(\frac{n\pi}{L\sigma} \right)^2 - 1} \right) \approx 0. \quad (\text{A.20})$$

We have made sure that this condition was met in the analyses presented. The criterion includes cases where the spatial extent of the correlations is relatively small compared to the size of the overall sensory sheet, which should be a realistic case in many biological scenarios. Finally, it should be pointed out that our procedure does not require the calculation of exact eigenvalues, but only relies on their relative ordering to be preserved and results should therefore be robust to small errors in the approximations.

D. Ordering in the 2D square case.

For rectangle regions, the ordering can be solved by calculating the number of lattice points enclosed by a quarter ellipse (Strauss, 2007). Here we use square regions and therefore the solution is the number of points enclosed in a quarter circle. The Gauss circle problem determines the number of integer lattice points which lie within a circle with radius $p \geq 0$, with its center at the origin:

$$N(p) = \#\{(l, m) \in \mathbb{R}^2 \mid l^2 + m^2 \leq p^2\}. \quad (\text{A.21})$$

The number of lattice points within the circle is approximately equal to its area. The number of points within a square region can be approximated by calculating the area of the upper quarter of the circle (positive values only).

$$N(p) = \frac{\pi p^2}{4} \quad (\text{A.22})$$

The number of eigenvalues in each region is therefore the area of the intersection of the circle and region.

For each region we calculate the number of lattice points enclosed by a quarter circle with radius equal to $l^2 + m^2$ for R_1 , and the solution of equation (14)– $n^2 + o^2$ – for R_2 , where $l, m, n, o = 1, 2, \dots$. This number is approximately the area of the quarter circle. For values of $l^2 + m^2$ or $n^2 + o^2$ greater than the total number of eigenvalues in each dimension ($L\sqrt{d}$), the approximation diverges from the true ordering as the area of the quarter circle becomes larger than the area of the lattice (region). In this case a correction term is added:

$$N(p) = \begin{cases} \frac{\pi p}{4} - p \arccos\left(\frac{k}{\sqrt{p}}\right) - k\sqrt{p-k^2}, & \text{if } \frac{k}{\sqrt{p}} < 1. \\ \frac{\pi p}{4}, & \text{otherwise.} \end{cases} \quad (\text{A.23})$$

where p is either $l^2 + m^2$ or $n^2 + o^2$ for R_1 and R_2 respectively, k is the total number of eigenvalues in each region. Assuming region size L , where each receptor is spaced one unit apart, $k = L^2$ for R_1 , and $k = (Ld)^2$ for R_2 .

E. Allocation for multiple regions in 2D.

For more than two regions, density and activation ratios for each additional region are calculated relative to a chosen baseline region. This leads to following general form for calculation of the eigenvalues of any Region x :

$$R_x: \lambda_{l,m} = \frac{2\sigma a \sqrt{d_x}}{l^2 \pi^2 L^{-2} \sqrt{d_b} + m^2 \pi^2 L^{-2} \sqrt{d_b} + \sigma^2 \sqrt{d_b}} \quad (\text{A.24})$$

where a is the region activation scaling ratio, d_b is the density of the baseline region, and d_x the density of region x . $l, m \in \mathbb{N}$ enumerate different eigenvalues for region x .

F. Allocation for the 1D case.

The 1D case for changes in density has previously been addressed in Edmondson et al. (2019). Here we extend this to include changes in activation. For two regions R_1 and R_2 , we can calculate their eigenvalues as:

$$R_1: \lambda_l^{(R1)} = \frac{2\sigma}{l^2 \pi^2 L^{-2} + \sigma^2} \quad (\text{A.25})$$

$$R_2: \lambda_m^{(R2)} = \frac{2\sigma ad}{m^2 \pi^2 L^{-2} + \sigma^2} \quad (\text{A.26})$$

where d is the ratio of higher and lower densities, a is the ratio of receptor activation, L is the length of the region, and $l, m \in \mathbb{N}$ denote successive eigenvalues for regions R_1 and R_2 , respectively.

To calculate how many output neurons are allocated to region R_2 as a function of the number of neurons allocated to region R_1 , we set $\lambda_l^{(R1)} = \lambda_m^{(R2)}$ and solve for m . This yields

$$m = \frac{\sqrt{ad(l^2 \pi^2 + L^2 \sigma^2) - L^2 \sigma^2}}{\pi}. \quad (\text{A.27})$$

It becomes apparent that for $l = 1$, i.e. the first neuron allocated to region R_1 , we have already assigned $m = \frac{\sqrt{ad(\pi^2 + L^2 \sigma^2) - L^2 \sigma^2}}{\pi}$ neurons to region R_2 . As we allocate more neurons to region R_1 , the ratio $\frac{m}{l}$ simplifies to: $\lim_{l \rightarrow \infty} \frac{m}{l} = \sqrt{ad}$. The fraction of neurons allocated to each region therefore depends on the size of the bottleneck and converges to $\frac{1}{1+\sqrt{ad}}$ and $\frac{\sqrt{ad}}{1+\sqrt{ad}}$ for R_1 and R_2 respectively.




Fig. S1. Effect of different values of σ on the covariance function decay. Smaller σ leads to a slower decay and therefore a larger extent of spatial correlations. When the sigma is large, receptors have less similar responses to their neighbours, and therefore co-vary less. Orange points denote the baseline spacing of receptors along one dimension. When receptor density is varied, the other region will have higher density (e.g. 1:3, indicated by blue dots).

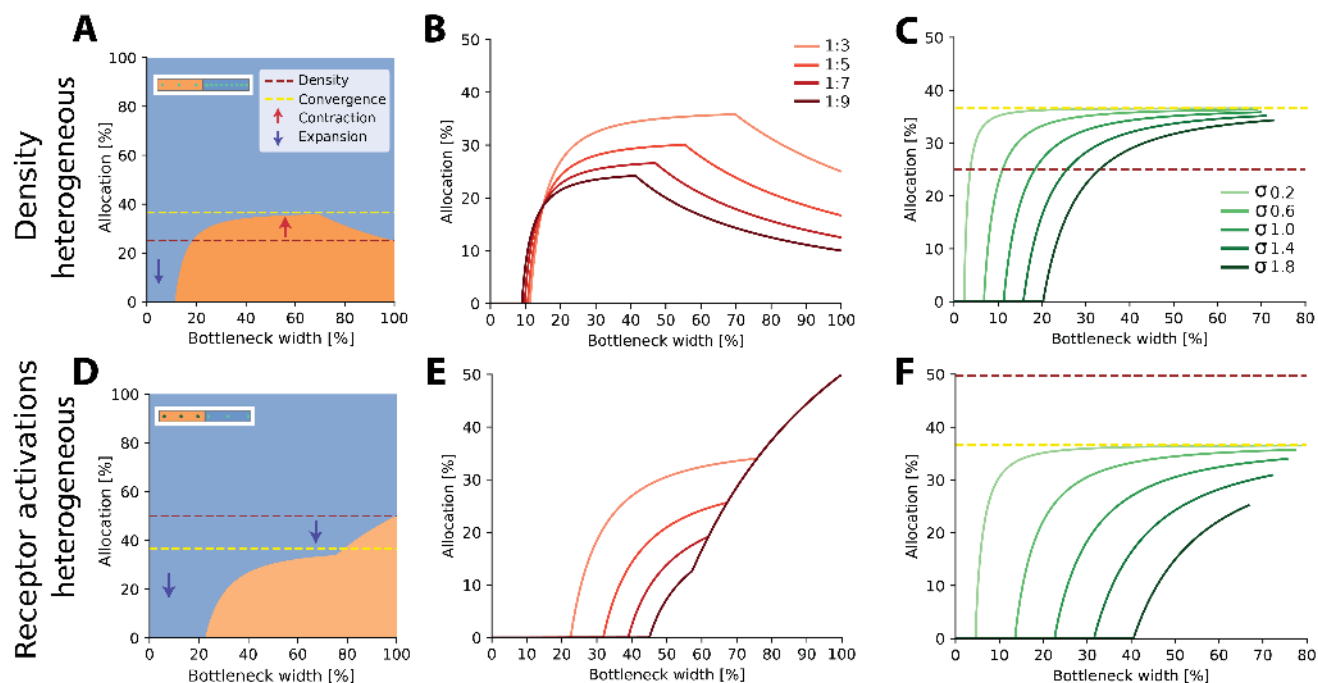


Fig. S2. Resource allocation for heterogeneous receptor densities and variations in input statistics in 1D. **A.** Examples of allocation for two input regions with differing receptor densities and activation (see insets) for different bottleneck widths, demonstrating complex trade-offs in resource allocation. Here ratio for low:high density = 1:3, $\sigma = 0.6$. Dotted lines show representation proportional to receptor density (red) and convergence in the limit (yellow). Arrows indicate expansion (down) and contraction (up) of the higher density (blue) region. Both expansion and contraction of the higher density region is seen, similar to the 2D case. **B.** Allocation boundaries for different density ratios. Ratios are given as low:high density. **C** Effect of changing the decay σ parameter. Density ratio is set at 1:3 for all sigmas. Panels **D-F**, same as above but for density constant across the two regions and receptor activation manipulated. Activation ratio is set at 1:3 in panels D and F. Note that in panel D, the representation of the higher activation (blue) region is expanded for all bottlenecks.

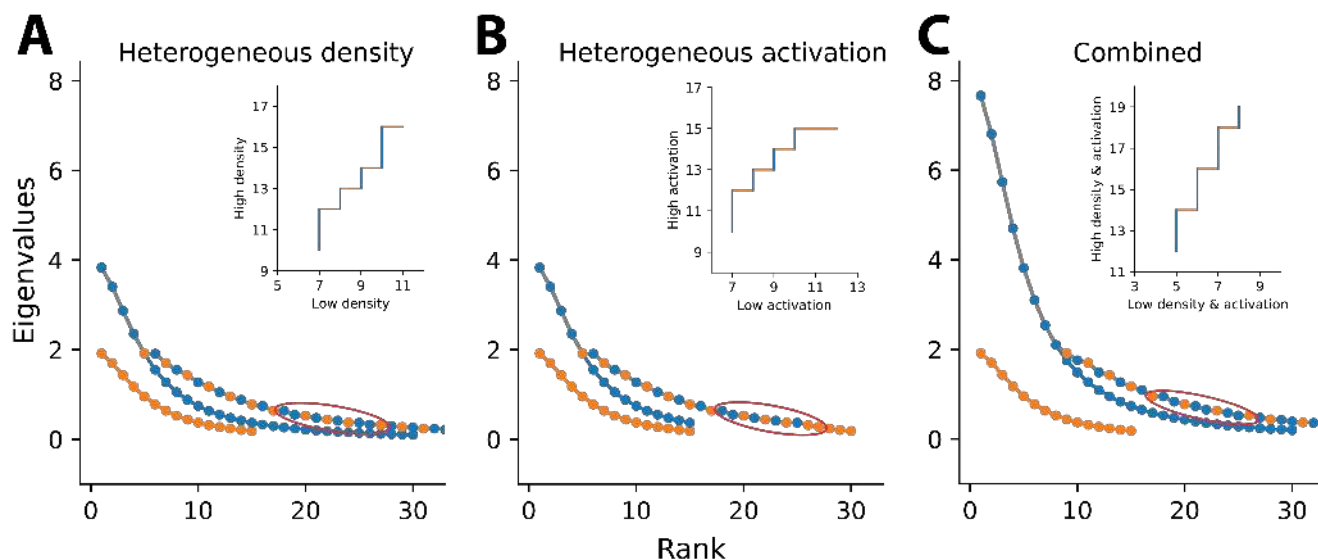


Fig. S3. Illustration of eigenvalue sorting and resulting allocation. Main panels: Eigenvalues for heterogeneous density (A), heterogeneous activation (B), or combined (C). For each the manipulated ratio is set as 1:2. For simplicity, the example considers 1D regions. Orange markers show eigenvalues from the baseline region, which is identical for all panels. Blue markers show eigenvalues from the second region with higher density, activation, or both, respectively. Markers connected by the grey line show the combined set of sorted eigenvalues from both regions. Insets: Cumulative allocation for the set of eigenvalues marked by the red ellipse in the main panel. Horizontal orange lines indicate eigenvalues from and therefore allocation to the orange region, whilst vertical blue lines indicate allocation to the blue region. **A.** Heterogeneous density. Increased receptor density in the blue region causes higher spatial correlations between neighbouring receptors and therefore larger eigenvalues, leading to a scaling of the eigenvalue curve. Increased density also leads to a larger number of receptors in the blue region, increasing the number of eigenvalues from this region and the total number of eigenvalues considered. **B.** Heterogeneous activation. Increased variance in the blue region leads to a scaling of the eigenvalue curve. However, the number of receptors in the blue region is not affected, explaining differences observed compared to A. **C.** Heterogeneous density and activation. Effects from A and B are combined, leading to a double-scaling of the eigenvalue curve, plus and increase in the total number of eigenvalues.

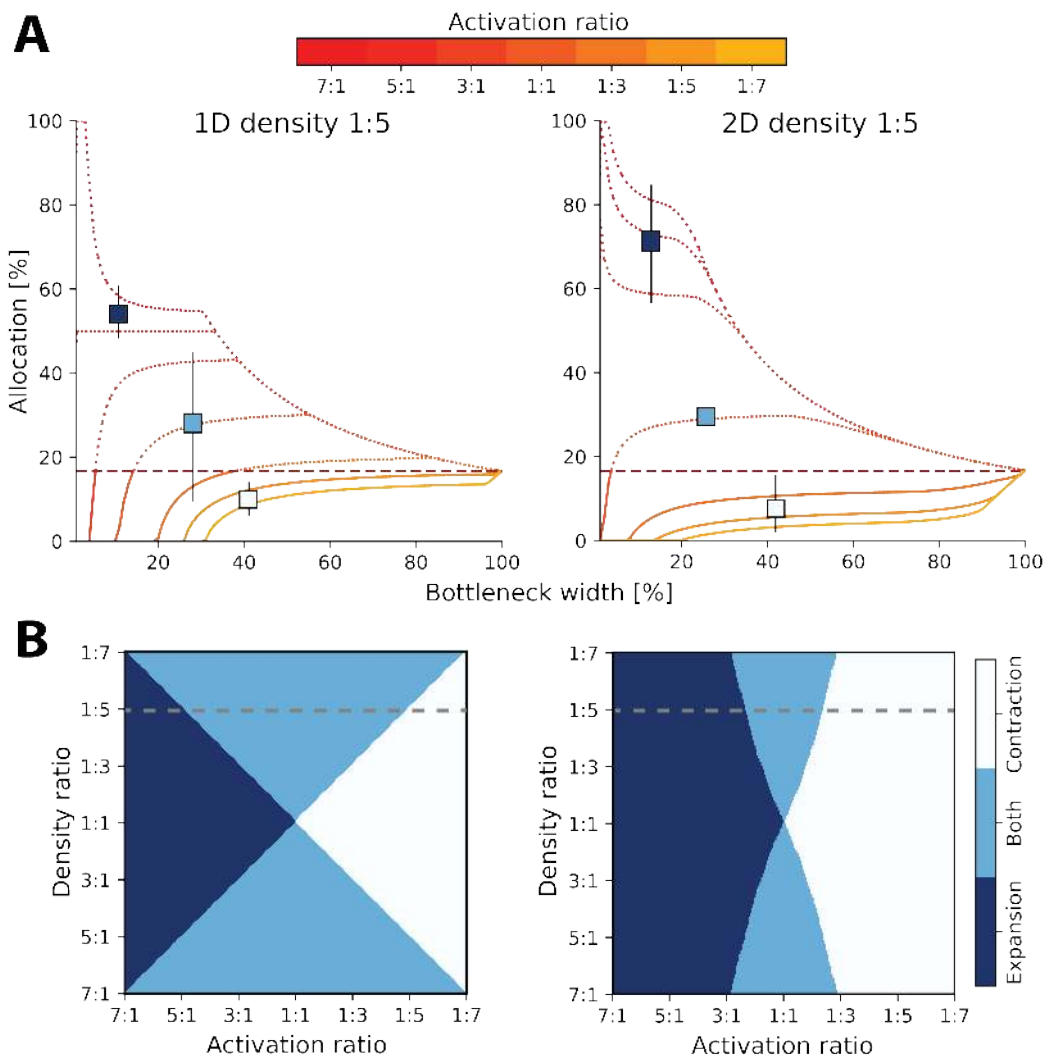


Fig. S4. Comparison between 1D and 2D results for heterogeneous activation and density. **A.** Effect of changing both the density and activation, and possible resource allocations for two regions. Plots show the same density ratio, 1:5, considering 1D (left) and 2D (right) regions. Allocation % is for baseline region (region 1). Each line denotes the variance ratios between Region 1:Region 2 (see colorbar). Lines which are solid denote contraction of region one's representation versus density, whilst dotted lines are expansion. The horizontal red dashed line shows the proportional density representation. Across both 1D and 2D, different patterns of expansion and contraction occur. Region one is more likely to either be fully expanded or contracted for the whole bottleneck in 2D, whereas in 1D both expansion and contraction occurs (denoted by light blue square, which covers three activation ratios in 1D, but only 1:1 in 2D as indicated by the vertical box lines). In 1D, the same activation and density ratio leads to convergence at the proportional density, whereas for 2D, the convergence is always lower. **B.** Comparison of the possible allocation schemes between 1D (left) and 2D (right). Dotted grey line denotes density ratio from panel A. In 2D, the likelihood of both expansion and contraction (light blue region) of the representations decreases.

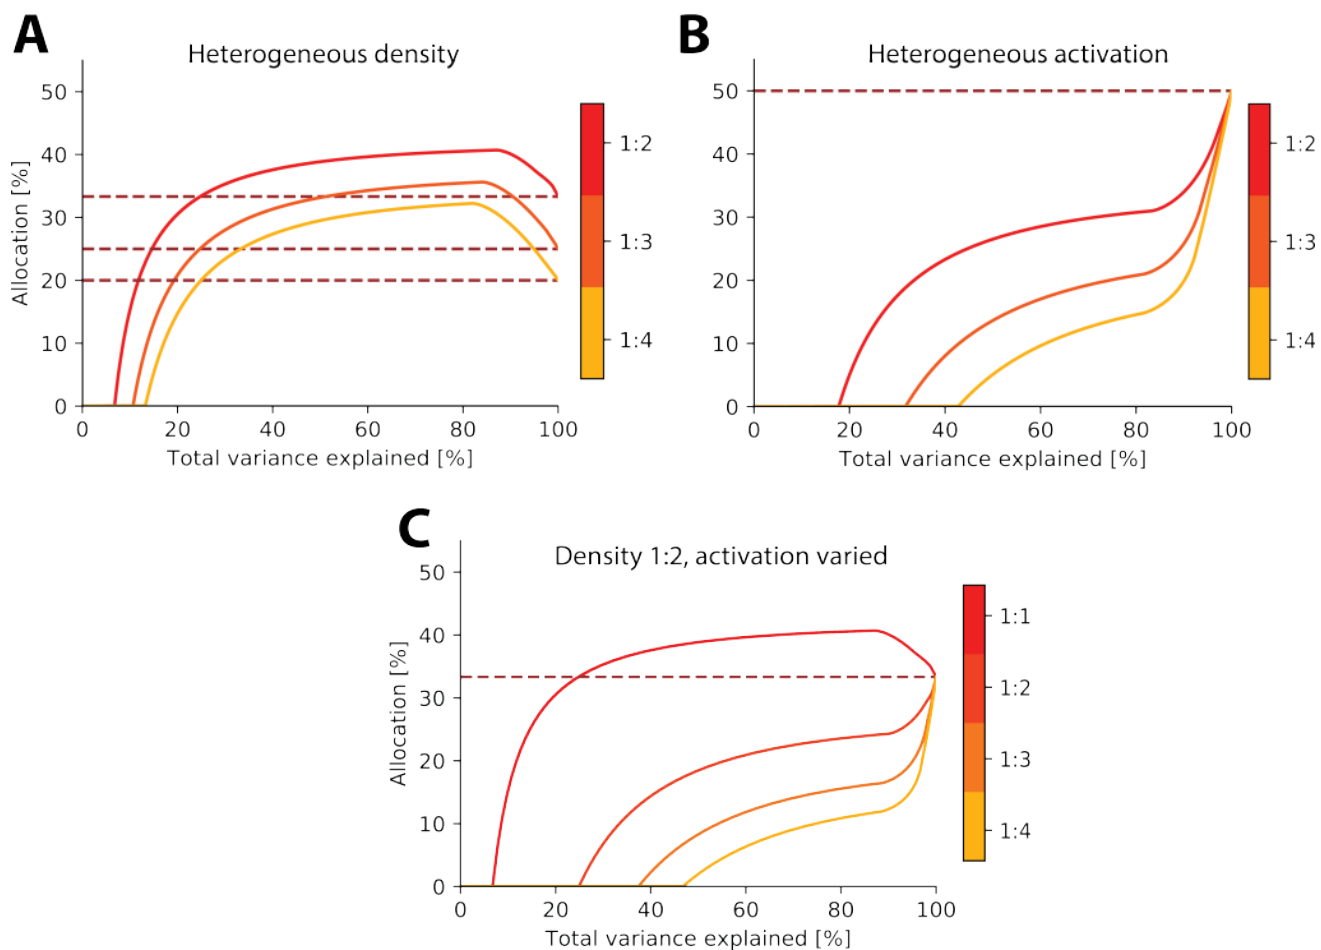


Fig. S5. Limit on information rather than number of neurons. Resource allocations where the bottleneck is expressed as variance explained. Because eigenvalues decrease dramatically in size, this re-expression results in a 'squashing' of the allocation curve towards higher bottleneck sizes, as the initial eigenvalues in the sorted set explain a much larger amount of the variance. **A.** Examples with heterogeneous density. **B.** Examples with heterogeneous activation. **C.** Examples with heterogeneous density and activation. Dashed red line in each denotes allocation proportional to density.